Requirements and Enforcement Points for Policies in Industrial Data Sharing Scenarios

Michael Falkenthal[1], Felix W. Baumann[2], Gerd Grünert[2],
Sebastian Hudert[2], Frank Leymann[1], Michael Zimmermann[1]

[1]Institute of Architecture of Application Systems,
University of Stuttgart, Germany,
[lastname]@iaas.uni-stuttgart.de

[2]TWT GmbH Science & Innovation,
Germany,
[firstname].[lastname]@twt-gmbh.de

BibTeX

```
@inproceedings{Falkenthal2017,
    author    = {Falkenthal, Michael and Baumann, Felix W. and Gr{\"{u}}nert,
                 Gerd and Hudert, Sebastian and Leymann, Frank and Zimmermann,
                 Michael},
    booktitle = {Proceedings of the 11\textsuperscript{th} Advanced Summer
                 School on Service Oriented Computing},
    pages     = {28--40},
    publisher = {IBM Research Division},
    title     = {{Requirements and Enforcement Points for Policies in Industrial
                 Data Sharing Scenarios}},
    year      = {2017}
}
```

Requirements and Enforcement Points for Policies in Industrial Data Sharing Scenarios

Michael Falkenthal[1], Felix W. Baumann[2], Gerd Grünert[2],
Sebastian Hudert[2], Frank Leymann[1], and Michael Zimmermann[1]

[1] University of Stuttgart,
Institute of Architecture of Application Systems
Universitätsstr. 38, 70569 Stuttgart, Germany
`[firstname].[lastname]@iaas.uni-stuttgart.de`
[2] TWT GmbH Science & Innovation
Industriestr. 6, 70565 Stuttgart, Germany
`[firstname].[lastname]@twt-gmbh.de`

Abstract. Industry 4.0 endeavours often integrate and analyze a multitude of data, such as data about machinery, production steps, and environmental conditions, in order to optimize manufacturing processes. Thereby, they aim to reveal information hidden in formerly isolated data silos via holistic analytics approaches. However, the integration of such data silos is often accompanied by challenges according legal regulations, organizational obstructions, and technical implementations, among others. Therefore, in this work we present a list of key challenges, which have to be commonly overcome in integration projects dealing with essential data from production processes. They can be used as a check list to address recurring challenges in future Industry 4.0 projects. Finally, we identify several plug-points in an abstract integration architecture, which have to be considered in concrete projects at hand to enforced the requirements.

Keywords: Requirements, Policies, Data Aggregation, Industrial Data, Data Integration, Industry 4.0

1 Introduction

The 4th industrial revolution, respectively known as Industry 4.0 [13], is facilitated by developments in the fields of data analytics, which evolve in this context to a new research field of so-called smart services [2]. Besides the availability of easily accessible cloud computing resources and advances in the miniaturization of sensors and Internet of Things (IoT) devices, the need for smarter factories and dynamic production processes are main drivers for manufacturing companies to foster new analytics approaches targeting the automated optimization of production lines. Thereby, different technologies, e.g., IoT, shall be leveraged in order to capture data about manufacturing environments and supporting processes in a very fine-grained manner. For instance, sensors are brought out into factories, which allow to enrich already present monitoring data about production processes with additional environmental parameters, such as temperature or humidity. The major goal of such endeavours is to identify previously hidden auxiliary conditions influencing production processes in a specific manufacturing facility. In ideal cases, conclusions about changing degrees of incorrectly produced parts can be deduced and machinery can be automatically adjusted appropriately to compensate such changing parameters. Another example is to integrate data silos from different production units to enable holistic analyses inferring new insights and optimization potentialities of production processes to make factories more adaptable to drastically increasing variations in the product portfolio.

However, while such developments promise to align whole industries for the upcoming era of dynamically and rapidly changing productions, actual Industry 4.0 projects are typically faced with different kinds of challenges that have to be considered in order to attain success. On the one hand, the integration of many different data sources and the analysis of big amounts of data both require immense expertise in terms of the development of analytics algorithms and the operation of integration middleware. On the other hand, the technical perspective of such an data analysis project is commonly not the most substantive one if it is about to prosper. Further requirements regarding law constraints, organizational obstacles, or qualities and semantics of analysis results have to be managed, which are often unapparent in the course of an Industry 4.0 project.

Therefore, we present findings by the research project *SePiA.Pro* [1, 16], which is located in the context of Industry 4.0. We describe and structure ascertained requirements concerning the integration and processing of business-critical data about manufacturing processes and production steps. Among the discussion of these requirements we also locate different enforcement points in an abstract and, thus, generic integration architecture that can be considered as technical hooks allowing to ensure compliance according the identified requirements.

The remainder of this paper is structured as following: we motivate the challenges of Industry 4.0 projects in more detail and give deeper background information in Section 2. We explain and structure key challenges in this context in Section 3 and discuss possible enforcement points to assure the compliant usage and security of business critical data in Section 4. Related work that supports and extends the presented findings in this work is discussed in Section 5. We conclude this work in Section 6 by summing up results that are expected from this work and the project and identify relevant future work.

2 Motivation and Background

Typical Industry 4.0 projects have to deal with the integration of formerly disconnected data silos, be it because of different production units, departments, or even legal entities. All these data sources are from the production or supply chain. Such a scenario is depicted in Figure 1, which illustrates the isolation of two production units on the right. In many companies, hierarchical organizations or the distribution of manufacturing among different production facilities lead to the emergence of data silos, i.e., technically disconnected databases. Each data

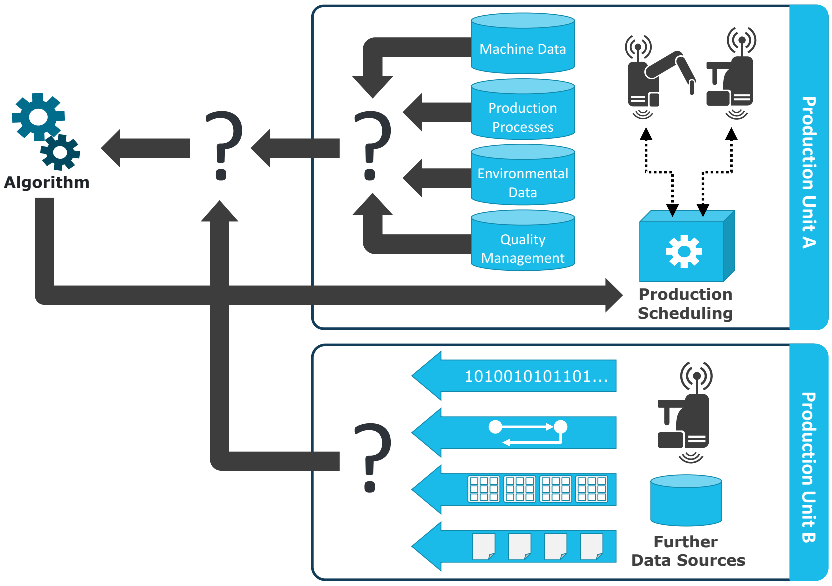

Fig. 1. Typical Industry 4.0 integration scenario: (i) different types of data are integrated as exemplarily depicted in *Production Unit A*, while (ii) different technical characteristics, such as data streams, data batches, and transmission via file are exemplarily illustrated in *Production Unit B*. Further, data from the different production units is further integrated and exposed to an external analytics *smart service*.

silo typically comprises different types of data, such as exemplarily illustrated in *Production Unit A* by machine-data, data about production steps and processes, data about environmental conditions, and data from quality management monitoring production lines. The combination of this data can enable the creation of additional knowledge that can benefit the data owner. Specific about this setup is that the different data types are usually disconnected even if they are technically available in a production unit. This is often due to different backend systems, diverging data formats and semantics, as well as missing adapters and integration middleware. Besides these technical impediments also organizational structures and responsibilities can cause such inhomogeneous system landscapes.

The situation is typically even more complex, since besides different types of data, also the technical characteristics of how it is captured and provided can differ greatly. This circumstance is illustrated in *Production Unit B*, where data is (i) available via *data streams*, i.e., continuous streams of bits and bytes, (ii) a *request – response* model, where an inquiring system sends specific requests for data, (iii) *batches of data*, such as results from executed sets of pooled SQL queries, or as (iv) a number of *file* exports. Therefore, also the heterogeneity if these characteristics adds to the complexity of integrating data from different systems. Thus, from a technological point of view, the integration of these data is a non-trivial task, as indicated by the question marks in the depicted production units in Figure 1. The integration is non-trivial because immense expertise about the source systems, as well as suitable integration middleware is required.

However, once data shall be integrated among different departments, production units, or even legal entities there often arise also obstacles in terms of legal issues and organizational caveats, which have to be handled by integration projects. Thereby, the integration of data from different legal units might cause violations of country-specific law, e.g., laws and acts intending to protect the economy from monopolistic consolidations of companies, consumer protection laws or laws against insider trading. This applies especially to integration scenarios, where competitors strive to cooperate in non-business-critical areas. Also, if data shall be integrated just internally, different departments or organizational units might inhibit the integration of their data, since analyzing data more holistically often allows for more transparencies and comparisons with other units, thus fueling rivalry. All these non-technical issues have to be clarified in order to provide proper boundary conditions for developing and applying algorithms to analyze connected data holistically, which is illustrated by the third question mark where inputs from the two production units are joined.

Based on this motivating scenario, a list of key challenges and requirements is described in Section 3. In combination with identified policy enforcement points in Section 4, they can be used as a starting point to identify, refine, and address such obstacles in future projects working in the context of Industry 4.0.

3 Requirements for Industrial Data Sharing Platforms and Policies

In this section, requirements for integration systems and data protecting policies, respectively, are described and categorized. The requirements can be classified into legal requirements, which are mainly dependent upon the geographic or legal location in which the data acquisition and processing is performed. The legal requirements are not entirely harmonized among different countries, making it important to involve legal experts early on in such a project.

Organizational requirements are usually defined and enforced by the company or companies employing such a smart service project. There is an overlap between organizational and regulatory requirements with regulatory requirements influencing and creating organizational requirements. Organizational requirements can, furthermore, be different within one single business entity, e.g., in a company which is operating in different countries. These requirements can further depend upon the customers affected by such smart services, e.g., the requirements can differ for corporate, private or governmental customers.

The third category are technical or structural requirements, which mainly stem from the underlying technical landscape. The fourth category comprises of the logical or principal requirements, and, therefore, groups those that are enforced to align with certain objectives or goals for which a particular smart service project is created. The requirements are identified in the following sections and described by respective examples and explanations.

3.1 Legal Requirements

In general, legal requirements arise from the country-specific legal situation. Thus, it is important to learn which laws have to be considered for elaborating a compliant data integration solution. This can be important in scheduling Industry 4.0 projects, because law dictates the fundamental frame about which data is allowed to be integrated. Thus, legal restrictions can prohibit to integrate data in order to be analyzed holistically, although there are no technical limitations.

The Data Privacy Act protects individual-related data in Germany. Data related to people is specifically protected to be processes arbitrarily. Transfered to data integration scenarios, e.g., data about bank accounts must not be combined with data about the purchasing behaviour of customers without being explicitly approved by them. Another example affects the processing of data about production processes. In such cases, it is often prohibited, or at least critical, e.g., to combine data about production processes, downtimes of conveyor belts, and staff for calculating the overall efficiency of personnel. Such purposes mostly have to be clarified with and approved by the employee organization of the company.

Sharing data between different legal entities or internal units can lead to legal issues. On the one hand, if data is shared with competitors in the same or in equal business areas, this can violate law against the suppression of competition, i.e., anti-cartel law. On the other hand, integrating data of distinct internal units can violate law and internal compliance policies if, e.g., the resulting integrated dataset and the automated analysis is not consistent with four-eyes-principles. Therefore, overcoming data silos might cause the circumvention of formerly established compliance processes because due to such regulations data must be distributed among separate units with different management responsibilities.

Anti-discriminatory algorithm design must be enforced in case of legal requirements to avoid biased algorithms. This can be of importance when algorithms and analyses involve person-attributable factors such as gender, religion or race. For industrial data scenarios this applies in cases when data about personnel, e.g., from processes, is combined with other data to draw conclusion's about their performance or qualification for specific tasks. Such analysis scenarios then typically have to deal with requirements detailed in *The Data Privacy Act* above.

3.2 Organizational Requirements

Organizational requirements typically stem from responsibilities of management staff, hierarchies, and segmentation of companies into departments, divisions, and units. Besides this, also sociological connections can influence data integrations among different departments. In any case, a integration project can benefit from a sponsor with wide-ranging responsibilities regarding these influencing factors.

New continuous transparencies of business unit data arise by enabling holisitc analyses. This is due to the fact, that formerly isolated data silos are connected and integrated, which can cause suspiciousness at any affected employee. Therefore, the intended transparencies as a result of the analysis of integrated data can, e.g., lead to unpleasant comparisons of the performance of different departments. Thus, people might refuse collaboration with integration projects because they fear being bad in comparison with others.

Data ownership is often dedicated to specific management responsibilities in a company. This implies that there are managers in a company, who must be enabled to enforce rules, i.e., policies about how specific data can be used and processed. However, data integration scenarios typically integrate datasets to allow comprehensive analysis. Thus, it has to be clarified how data policies can be enforced in the integrated dataset and if new data responsibilities are added.

Inter-company analytics scenarios are often not the major aim of Industry 4.0 projects. Nevertheless, abstracting and aggregating data until they no longer contain business critical information can still open up analysis scenarios along with further companies, which can lead to overall results and value adds for all participants. For instance, if companies operating machines share data about their processing environments, environmental conditions, and machine parameters with the machine vendors, this can enable completely new business models. Of course, then data has to be shared in a way that no business critical information, such as information about the produced parts, is captured. However, this can then enable the analysis of overall machine fleets by machine vendors resulting in suggestions about how to optimize the operation of machines.

3.3 Technical or Structural Requirements

The technical or structural requirements involve issues and obstacles, which can occur due to implications based on technological restrictions or implementation-specific difficulties. These restrictions typically have to be managed in the development and implementation phase of an integration project, while requirements as presented in the sections above have to be carefully considered and incorporated.

Different semantics of data from different data sources can lead to immense integration efforts. For instance, machinery from different vendors can be technically integrated based on compatible protocols but usually provide a vendor-specific data model. So, the different data models have to be compared and mapped to each other in order to assure precise semantics of the resulting integrated set of data. Often, a normalized data model has to be derived and additional data transformations have to be introduced for source systems to match with the integrated data model. Such transformations typically lead to more complexity in the overall integration system due to additional processing components but also because of additionally required processing infrastructure.

Different formats, quality of data, and acquisition rates are also obstacles, which have to be managed in order to enable the holisitc analysis of different data sources. Thereby, semantically equal values have to be adjusted, e.g., the fractual part of floating-point numbers. Another problem arises, if data with different accuracies is collected. Some data sources can provide a higher degree of uncertainity with some data than others, which must be reflected for decisions based on them. Finally, sensors and other data sources often provide data by different rates, which has to be considered by normalizing such data streams.

Data policies have to be inseperable from the data to be protected and integration middleware has to enforce them. This assures, especially in the field of industrial data, the required degree of security for business critical data. For instance, if a specific set of data is classified, i.e., it is defined that it must not leave the company, this has to be attached to the data in the form of a data policy, which can be processed and enforced by integration systems [15].

Arbitrary data transfers have to be secured by data policies. Hence, data policies have to be attached to data independently from the communication channel, be it the transfer of data via data streams, batch jobs, or ordinary files.

Policies have to be combinable on all aggregation steps as motivated in Section 2. Integrating data among different departments, business units, or even companies often implies that data is integrated on a cascade of different integration systems. Each integration step can require to initiate aggregations and obfuscation of data in order to enforce attached policies. However, in such scenarios policies also have to be applied in combination, i.e., policy aggregations have to be conceptually possible and must also technically be enforced.

3.4 Logical or Principal Requirements

The requirements presented in this section add general aspects to the above presented ones. They influence the quality of analysis results and the protection of data by adding general properties to be incorporated into integration systems.

Data results must yield a specific format or data-range. Thus, specific expectation checks should be applied to data at the different integration and aggregation steps as presented above. For instance, data input and data output at a particular processing step must conform to specific rules, which have to be defined. This assures that data does not get corrupted during different manipulations.

Enforceability of data policies has to be assured twofold: firstly, policies must be enforceable under specific conditions, such as in accordance with time. So the relevance of a policy can be restricted via time constraints in a way, that it only applies for data with specific time stamps. Secondly, policy checking and enforcement must be automated to assure performance of integration scenarios.

4 Policy Enforcement Points in Industrial Settings

Policies must be enforced in any system. With the proposed and described system a distributed data network is created. By the nature of this distribution, different logical points are possible and reasonable for enforcement operations. By enforcing policies in different locations, different results and implications are manifested. In the current scenario, data is acquired and handled within *DataHubs*, software components that are equipped with control and access logic, thus, managing the access and acquisition of data for the stakeholders. They can be recursively stacked within enterprises and locations as the following example depicts.

A DataHub is placed logically near the data producing machine and unifies the access to this device such that it can be used in the resulting smart service. Another DataHub is placed within one factory building, aggregating data from and unifying access to several downstream DataHubs located at various machines. Furthermore, a DataHub is placed at a business unit that controls various factory buildings and individual machines with associated DataHubs. The location and placement of the DataHubs can be categorized as follows:

- Directly at the data-producing machine. This placement requires knowledge about access control structures, based on employees or groups, which might not be available in this level as the management is usually a few layers up.
- Aggregating data within a physical location, e.g., within a factory building.
- Within a business unit, responsible for data acquisition of a various number of physical and logical locations.

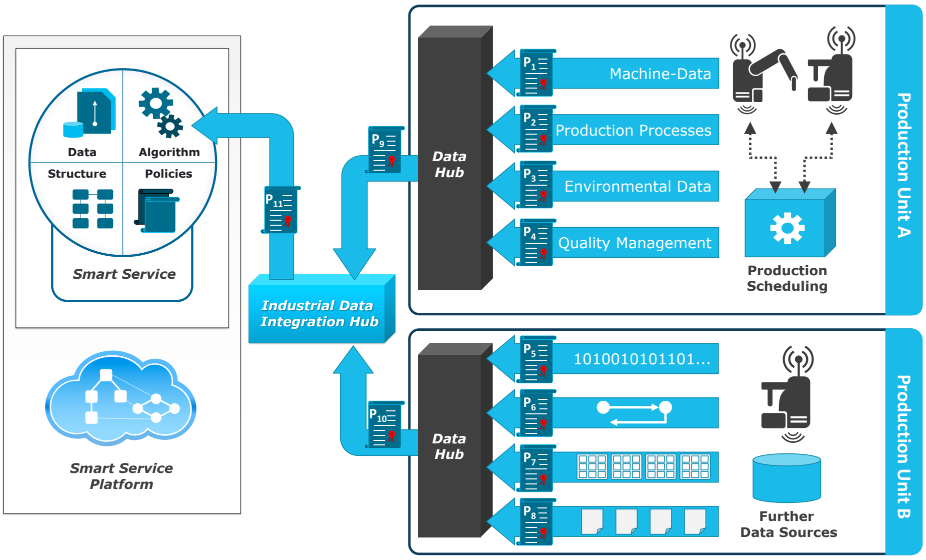

Fig. 2. Data policy enforcement points along different aggregation steps as indicated by eleven policy documents in an abstract integration scenario.

- At company borders, for inter-company connection of data processing tasks. This placement requires, like all others too, detailed knowledge of the connected machinery to enforce certain access restrictions on specific data fields, which might be not available this far from the data production.

All placements vary in the degree of implications caused. To summarize it can be stated, that for the data access enforcement, knowledge about the allowed and disallowed groups and persons, knowledge about company structure, aims, and targets, and, finally, knowledge about the data structure available is required. The different involved parties have varying knowledge about this.

The DataHub is intended to be data-agnostic, meaning that the access is unified, but for the implementation this knowledge is still required. The individual DataHubs do not have any knowledge of the structure further upstream and are only aware of their immediate downstream level. The logic is imbued into the system through the smart services, which have a complete picture of the connectedness of the machines and DataHubs. The implementation scenario of the DataHubs and the machinery within the *SePiA.Pro* project is depicted in Figure 2. In this figure, the *Industrial Data Integration Hub* is also a DataHub component, that is named to indicate its purpose.

The policy enforcement can further be divided by being upstream or being downstream. Upstream denotes the case, where all data access is propagated from the user and the data is acquired accordingly and only in the situation of the transmission of the data back to the user, it is checked if the data is allowed to be procured to the user. The downstream policy checking works by testing the data acquisition or processing requests prior to their execution and acquisition. Aggregation of data might result in new data that the user is not allowed to process which is possible to filter in the upstream processing. In the downstream processing, the user might be wrongfully inhibited to generate data requests that would, eventually, through aggregation or pre-processing, result in data the user would be allowed to process. As an example, a policy could prohibit the usage of personal information, such as which operator of a machine is working with a machine for how long. In this example and with downstream policy application, it is possible to query the machine operating hours and separately shift plans for the respective machine. Both these queries do not yield the forbidden information, so it would be allowed. Further combination of these data would yield the forbidden information but could not be prevented in this case. To prevent this, upstream policy application is required.

The policies, indicated by the scroll symbol and named P_x in Figure 2, can be attached to raw data and processed data of various types. The policies are ensured to be enforced by the DataHub. In the reverse direction, sending instructions towards the data sources, e.g., for the addition of additional sensors or for restructuring of the data, are also possible with the system. This direction is not depicted in the figure. These instructions are also enhanceable with policies as described.

5 Related Work

In the following, related work is discussed, which extends the context of this work. Thereby, we especially point out work presenting details and approaches about the presented challenges and starting points to solve them in particular projects.

According to Weyer et al. [24], one additionally identified and important challenge for the advent and success of smart factories is the standardization of protocols, technologies, and data formats. They identify that production systems, nowadays, are still only vendor-specific ecosystems, which are not driven by open standards. This hinders their interplay and integration to foster automatic control and adjustment of production processes. Thus, they provide *SmartFactory*^{KL} as an exemplary reference for a modular and adaptable production system.

Regarding the communication and connectivity of machinery, devices, and further data sources, Varghese and Tandur [23] describe the key role of wireless communication networks to enable Industry 4.0 systems. They discuss current key challenges in the field of wireless communication and argue how the 5th generation of wireless networks can tackle these. The identified key challenges concerning wireless communication extend the presented technical challenges in this work. Wollschläger et al. [25] further underline these challenges by identifying IoT as the leading technological evolution, which enables the development of smart platforms to access and orchestrate industrial data and devices.

To overcome the technical challenges in terms of the design and implementation of integration architectures as abstractly depicted in this work, there are different approaches presented. General IoT reference architectures to identify main system components are given, such as presented by Guth et al. [12]. The enterprise integration patterns by Hohpe and Woolf [14] provide best practices to design and implement integration systems. Implementations of these patterns are already available via different middleware technologies, such as Apache ActiveMQ [21], Apache Camel [22] or Spring Integration [17]. Further, the cloud computing patterns by Fehling et al. [11] provide knowledge about integration of private and public clouds, which can get necessary in Industry 4.0 endeavours if local processing power is not sufficient to execute analytics algorithms contained in smart services. To specifically deal with architectures and characteristics of IoT-related systems and devices, Reinfurt et al. [18, 19] provide a collection of *Internet of Things Patterns*, which they plan to develop towards a pattern language for IoT. The presented requirements in this work can be mapped to their patterns in order to find proper solution concepts. Finally, to ease and guide design and implementation of IoT systems and integration scenarios and, thus, to efficiently overcome the identified implementation challenges in this work, Falkenthal et al. [5, 4, 6, 9] describe approaches to connect concrete implementations to patters using pattern and solution repositories as introduced by Fehling et al. [10].

Finally, technologies from the domain of cloud computing have been identified to be drivers of the 4th industrial revolution in terms of automating the provisioning and management of analytics stacks [8] and to enable function and data shipping scenarios based on situational conditions [7], such as legal and organizational requirements as identified in this work. Application of cloud technology is also considered essential in manufacturing concepts, such as cloud manufacturing as described in Baumann et al. [3], where the connection of additive manufacturing technology to the Internet is described, thus enabling collaborative work.

6 Conclusion and Future Work

We presented in this work findings from the project SePiA.Pro [1], which investigates the issues and challenges of Industry 4.0 projects in terms of the research program Smart Service World of the federal ministry of economics and energy of Germany. Thereby, we elaborated requirements for protecting industrial data in the context of Industry 4.0 endeavours via requirements or data policies, respectively. Such data policies are means to specify constraints, restrictions or instructions that apply to the data, taking into account aspects such as data accessibility, utilisation, processing, obfuscation, storage or generation. The policies extend common access control rules and restrictions to incorporate concepts such as temporal, logical and organisational triggers. An exemplarily scenario for enabling trust and enforcing implementations was analysed within this work, which can be used as a coarse-grained overview to attach data policies to relevant data sources and plug-points in data integration architectures. The rationale for such an attachment of policies is to secure and protect data from manufacturing environments in standards-based deployment models such as cloud computing. These models can be used to provision smart services and wiring them with arbitrary data sources, such as databases, data aggregation services, industry specific machine to machine or IoT related data streaming endpoints.

In future work, we plan to further investigate, how and through which means, i.e. systems and parties, the identified challenges, requirements, and policies can be enforced at several points in time of the lifecycle of smart services — specifically at modelling time, deployment time and runtime — to overcome the above mentioned obstacles. Based on these investigations we plan to extend the open-source provisioning engine OpenTOSCA to enable the enforcement of data policies in Industry 4.0 deployment scenarios as presented by Falkenthal et al. [8] and also or more general IoT integrations such as presented by [20].

Acknowledgments. This work is partially funded by the project SePiA.Pro (01MD16013F) of the BMWi program Smart Service World.

References

1. Service Plattform for the Inteligently Optimization of Manufacturing Environments, `http://projekt-sepiapro.de/en`, last accessed on 28th July 2017
2. Allmendinger, G., Lombreglia, R.: Four strategies for the age of smart services. Harvard Business Review 83(10), 131 (2005)
3. Baumann, F.W., Eichhoff, J., Roller, D.: Collaborative cloud printing service. In: Cooperative Design, Visualization, and Engineering - 13th International Conference, CDVE 2016, Sydney, NSW, Australia, October 24–27, 2016, Proceedings. pp. 77–85. Springer (Oct 2016)
4. Falkenthal, M., Barzen, J., Breitenbücher, U., Fehling, C., Leymann, F.: Efficient Pattern Application: Validating the Concept of Solution Implementations in Different Domains. International Journal On Advances in Software 7(3&4), 710–726 (Dec 2014)
5. Falkenthal, M., Barzen, J., Breitenbücher, U., Fehling, C., Leymann, F.: From Pattern Languages to Solution Implementations. In: Proceedings of the Sixth International Conferences on Pervasive Patterns and Applications (PATTERNS 2014). pp. 12–21. Xpert Publishing Services (May 2014)
6. Falkenthal, M., Barzen, J., Breitenbücher, U., Fehling, C., Leymann, F., Hadjakos, A., Hentschel, F., Schulze, H.: Leveraging Pattern Application via Pattern Refinement. In: Proceedings of the International Conference on Pursuit of Pattern Languages for Societal Change (PURPLSOC 2015). epubli (Jun 2015)
7. Falkenthal, M., Breitenbücher, U., Christ, M., Endres, C., Kempa-Liehr, A.W., Leymann, F., Zimmermann, M.: Towards Function and Data Shipping in Manufacturing Environments: How Cloud Technologies leverage the 4th Industrial Revolution. In: Proceedings of the 10th Advanced Summer School on Service Oriented Computing. pp. 16–25. IBM Research Report, IBM Research Report (Sep 2016)
8. Falkenthal, M., Breitenbücher, U., Képes, K., Leymann, F., Zimmermann, M., Christ, M., Neuffer, J., Braun, N., Kempa-Liehr, A.W.: OpenTOSCA for the 4th Industrial Revolution: Automating the Provisioning of Analytics Tools Based on Apache Flink. In: Proceedings of the 6th International Conference on the Internet of Things. pp. 179–180. IoT'16, ACM (2016)
9. Falkenthal, M., Leymann, F.: Easing Pattern Application by Means of Solution Languages. In: Proceedings of the 9th International Conference on Pervasive Patterns and Applications. pp. 58–64. Xpert Publishing Services (2017)
10. Fehling, C., Barzen, J., Falkenthal, M., Leymann, F.: PatternPedia – Collaborative Pattern Identification and Authoring. In: Proceedings of PURPLSOC (Pursuit of Pattern Languages for Societal Change). The Workshop 2014. pp. 252–284 (Aug 2015)
11. Fehling, C., Leymann, F., Retter, R., Schupeck, W., Arbitter, P.: Cloud Computing Patterns: Fundamentals to Design, Build, and Manage Cloud Applications. Springer (Jan 2014)
12. Guth, J., Breitenbücher, U., Falkenthal, M., Leymann, F., Reinfurt, L.: Comparison of iot platform architectures: A field study based on a reference architecture. In: 2016 Cloudification of the Internet of Things. IEEE (2016)
13. Hermann, M., Pentek, T., Otto, B.: Design Principles for Industrie 4.0 Scenarios. In: Proceedings of the 49th Hawaii International Conference on System Sciences (HICSS). pp. 3928–3937 (2016)
14. Hohpe, G., Woolf, B.: Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Solutions. Addison-Wesley (2004)

15. Pearson, S., Casassa-Mont, M.: Sticky policies: An approach for managing privacy across multiple parties. Computer 44(9), 60–68 (Sept 2011)
16. Pfeil, M., Odefey, U., Ritter, Y., Schuermann, M., Fäßler, V.: Smart services - the smart implementation of Industry 4.0. In: NAFEMS Seminar Simulation von Composites – Bereit für Industrie 4.0? (10 2016), https://www.researchgate.net/publication/317687064_Smart_services_-_the_smart_implementation_of_Industry_40, last accessed on 28th July 2017
17. Pivotal Software: Spring Integration, https://spring.io/spring-integration, last accessed on 28th July 2017
18. Reinfurt, L., Breitenbücher, U., Falkenthal, M., Leymann, F., Riegg, A.: Internet of things patterns. In: Proceedings of the 21st European Conference on Pattern Languages of Programs. ACM (2016)
19. Reinfurt, L., Breitenbücher, U., Falkenthal, M., Leymann, F., Riegg, A.: Internet of things patterns for devices. In: Ninth international Conferences on Pervasive Patterns and Applications. pp. 117–126. Xpert Publishing Services (2017)
20. da Silva, A.C.F., Breitenbücher, U., Hirmer, P., Képes, K., Kopp, O., Leymann, F., Mitschang, B., Steinke, R.: Internet of Things Out of the Box: Using TOSCA for Automating the Deployment of IoT Environments. In: Proceedings of the 7th International Conference on Cloud Computing and Services Science (CLOSER). pp. 358–367. SciTePress Digital Library (Jun 2017)
21. The Apache Software Foundation: Apache ActiveMQ, https://activemq.apache.org, last accessed on 28th July 2017
22. The Apache Software Foundation: Apache Camel, http://camel.apache.org, last accessed on 28th July 2017
23. Varghese, A., Tandur, D.: Wireless requirements and challenges in industry 4.0. In: 2014 International Conference on Contemporary Computing and Informatics. pp. 634–638 (Nov 2014)
24. Weyer, S., Schmitt, M., Ohmer, M., Gorecky, D.: Towards industry 4.0 - standardization as the crucial challenge for highly modular, multi-vendor production systems. IFAC Symposium on Information Control Problems in Manufacturing-PapersOnLine 48(3), 579 – 584 (2015)
25. Wollschlaeger, M., Sauter, T., Jasperneite, J.: The future of industrial communication: Automation networks in the era of the internet of things and industry 4.0. IEEE Industrial Electronics Magazine 11(1), 17–27 (March 2017)